April 2007

Contact	Phone
New York	
Richard Cantor	1.212.553.1653
David Hamilton	
Jennifer Tennant	

Confidence Intervals for Corporate Default Rates

Summary

- This Special Comment presents confidence intervals around historical average cumulative default rates by rating category over multiple horizons. The confidence intervals are derived by "bootstrapping" (creating 10,000 hypothetical data sets from) the original data set of 11,370 corporate rating histories by sampling with replacement.
- The results indicate that historical mean speculative-grade default rates are generally measured fairly precisely, with standard errors less the 10% of the estimated means. Investment-grade default rates, however, are measured much less precisely, particularly for issuers rated single A or above. Precision increases at longer horizons. (See the table below.)
- Moody's long-term ratings satisfy the Basel II criteria for effectively distinguishing relative credit risk. This is true even for "low-default portfolio" portion of the rating scale — letter ratings Aaa, Aa, and single A. The Aaa and Aa default rates are statistically differentiated at the three-year investment horizon, and the Aa and single-A default rates are differentiated at all horizons greater than one year.
- Such confidence intervals measure uncertainty about long-run means. In contrast, the uncertainty about the default rate of any individual ratings cohort is much greater, as revealed by the time series variation in default rates of cohorts formed at different monthly intervals.
- Although the historical variation in cohort default rates has been great, Moody's broad letter ratings consistently rank order default risk within cohorts. For example, since 1992, no higher rated cohort has experienced a higher 5-year CDR than any lower-rated cohort formed on the same month.

Variability of Cumulative Default Rate Estimates

	----- 1-Year Default Rates-----		--- 3-Year Cumulative Default Rates ---		--- 5-Year Cumulative Default Rates ---	
	Mean	**Stdev**	**Mean**	**Stdev**	**Mean**	**Stdev**
Aaa	0.00%	0.00%	0.00%	0.00%	0.10%	0.07%
Aa	0.01%	0.01%	0.04%	0.02%	0.18%	0.06%
A	0.02%	0.01%	0.22%	0.04%	0.47%	0.08%
Baa	0.18%	0.03%	0.93%	0.10%	1.94%	0.19%
Ba	1.20%	0.08%	5.57%	0.30%	10.21%	0.53%
B	5.23%	0.18%	17.04%	0.56%	26.79%	0.90%
Caa-C	19.47%	0.75%	39.73%	1.59%	52.66%	2.33%

Exhibit 16

Introduction

Moody's historical average cumulative default rates by rating category and investment horizon are among its most widely referenced statistics. In any finite sample, however, the historical average (or mean) default rate may overstate or understate the underlying population's true risk of default, depending upon whether the particular set of issuers included in the sample happen to experience lower or higher than expected default incidence. Quantitative credit analysts are, therefore, not only interested in mean default rates; they are also interested in their standard errors, which they can use to form confidence intervals around the estimated means.

In addition, banks that map their internal rating systems to Moody's rating scale should be interested in the confidence intervals around the average default rates by rating category. In particular, banks may want to know whether the default rates associated with the so-called "low default portfolio" portion of the rating scale (single A, Aa, and Aaa credits) are sufficiently differentiated by rating category to meet Basel II criteria for internal ratings systems.

In this Special Comment, we estimate confidence intervals for Moody's average default rates by rating category and investment horizon, based on corporate issuer rating histories from 1970-2006. The work extends research in the academic literature [Hanson and Schuermann (2006)] on one-year default rates to the multi-year horizon case.

We find that for the estimated mean default rates associated with the Baa and the speculative-grade rating categories, the standard errors are generally small, at 10% or less than the mean. The mean defaults for the Aaa, Aa, and A rating categories, however, tend to be less precisely estimated, with standard errors ranging from 15% to about 100% of the means, depending on the rating category and the investment horizon. We also find that the precision of the estimate of the mean tend to increase with the length of the investment horizon, because the mean default rate rises faster than the standard error as the horizon increases.

The estimated confidence intervals suggest that Moody's long-term ratings satisfy the Basel II criteria for effectively distinguishing relative credit risk. This is true even for the "low-default portfolio" portion of the rating scale — letter ratings Aaa, Aa, and single A. The Aaa and Aa default rates are statistically differentiated at the three-year investment horizon, and the Aa and single-A default rates are differentiated at all horizons greater than one year.[1]

As a byproduct of this analysis, we also present confidence intervals one-year hazard rates, which are also known as marginal default rates. The marginal default rate of year t is the probability of default in the year t, conditional on not defaulting during the prior t-1 years. Interestingly, as t increases and the horizon is lengthened, the precision of the estimated mean marginal default rates (as measured by the inverse of the coefficient of variation) improves more slowly than does the cumulative default rate and even decreases in many cases.

It should be noted, however, that standard errors around estimated long-run average default rates should not be confused with the much greater bands of uncertainty associated with expected performance of particular cohorts of issuers formed at specific points in time. Long-term default rates average performance across many macroeconomic scenarios and credit market environments; whereas, the default rate associated with any particular cohort will depend upon the realization of a particular macroeconomic scenario and credit market environment.

1. *At the one-year horizon, the default rates associated with these rating categories are properly rank ordered, but their differences may not be significantly different in the statistical sense. At longer horizons, however, their differences are clearly significantly different. These long-term horizon differences can be used to validate the effectiveness of these rating categories in distinguishing risk for internal rating systems. Hence, they can be used to differentiate risk within internal rating systems. [See Basel Committee (2005), page 5.]*

Exhibit 16

Methods for Measuring Long-Run Average Default Rates and Their Confidence Intervals

The academic literature suggests multiple approaches — analytic, parametric, and non-parametric (i.e., bootstrap) methods — for estimating confidence intervals around historical long-run average one-year default rate estimates.[2]

The simplest possible analytic approach assumes that default events are independent, and whether or not an issuer defaults is determined by a binomial process governed by a single known parameter (the probability of default). In large samples, the probability distribution of the empirical average default rate is known to be normally distributed with a mean equal to the population's underlying default rate and a standard deviation that is a decreasing function of the size of the sample. The utility of this statistic in measuring one-year default rate confidence intervals is discussed in Cantor and Falkenstein (2001), Stein (2006), and Hanson and Schuermann (2006). Cantor and Falkenstein (2001) also discuss how the approach can be adapted to take into account common temporal shocks to the aggregate default rate, which is equivalent to the use of the correlated binomial distribution [see Witt (2004)].

Another approach is to specify and to estimate parameters governing the evolution of the default process over time. Such parametric methods (by adding structure) use available data very efficiently; however, they can also be misleading if they are based on inaccurate models of the default process. In fact, in common applications, most parametric models appear to be based on assumptions directly at odds with reality and therefore likely lead to invalid estimates.[3]

As a practical matter, even if parametric and analytic approaches were to prove useful for the analysis of short-term default rates, they would remain difficult to adapt to the analysis of multi-year cumulative default rates. In order to calculate long-term default rates from a mix of short-term and long-term rating histories, Moody's (but not all other rating agencies) follows the standard academic practice of adjusting its default rates for the censoring introduced by the short-lived rating histories, which if left unadjusted lead to artificially low estimates of long-term default rates. As discussed in Hamilton and Cantor (2006), Moody's calculates long-term default rates by compounding one-year marginal default rates derived from cohort data constructed at monthly intervals. This method maximizes the use of the information available in the data set; however, it implies that the sample size for a multi-year default rate estimate is not easily defined, since the multi-year default rate is derived from a sequence of one-year marginal default rates, each with its own sample size.

In contrast, a non-parametric bootstrap method makes no assumptions about the statistical process generating the data, except that each issuer's rating and default history are independent.[4] Standard error estimates and other statistical inferences can be made based on Monte Carlo sampling using the actual distribution of the historical data. Bootstrap methods are equally appropriate for estimating confidence intervals for one-year and for multi-year cumulative default rates.

2. *For a summary and analysis of methods for estimating confidence intervals for the long-term average one-year default rate, see Hanson and Schuermann (2006).*
3. *For example, Hanson and Schuermann (2006) show that parametric continuous time estimators imply much lower default rates than observed in practice. This result is unsurprising since commonly used continuous time estimators assume rating transitions are Markovian, an assumption that is inconsistent with the rating change momentum that is strongly evident in the data [see Hamilton and Cantor (2004)]. The advantages and practical limitations of deriving non-Markovian continuous time estimators are explored by Christensen, Hansen, and Lando (2004).*
4. *The assumption of the independence of issuer rating histories is, strictly speaking, violated in reality. However, the assumption minimizes any correlation in default times that may exist in the data. For example, the default time for an issuer first rated in, say, 1974 is unlikely to be strongly correlated with the default time for an issuer whose rating history begins in 1998.*

Bootstrap Estimates of Multi-Period Default Rates and Their Confidence Intervals

Our bootstrap analysis follows closely the approach taken in Hanson and Schuermann (2006), except we extend their analysis of one-year default rates to include multi-year CDRs.[5] We create 10,000 data sets by sampling with replacement from our original historical data set of 11,370 firm rating and default histories, so that each of the 10,000 data sets itself has 11,370 firm rating and default histories. As a result, a firm which naturally appears only once in the original data set, will on average not appear at all in 5,000 of these data sets. However, the same firm's history will appear multiple times in many of the other data sets, so that the total number of times it will appear across all the data sets will average 10,000.

The bootstrapped data sets can be used to construct 10,000 observations on any particular sample statistic. For example, Exhibit 1 depicts the frequency distribution of all 10,000 of the estimated 5-year cumulative default rates for issuers rated Baa on each cohort's formation date. Notice that the mean for this sample is 1.94%, which is the same as that reported for the 5-year Baa default rate in our annual default study. The distribution around the mean, however, reveals that resampling 10,000 times from the original dataset could conceivably have produced a default rate as low as 1.35% and as high as 2.50%. Yet both of these outcomes are extraordinarily unlikely. The 5% and 95% confidence level points are more meaningful measures of the range of potential outcomes we are ever likely to see. As revealed in the chart, 95% of the time, we would expect to see a long-term average 5-year Baa default rate at least equal to 1.63% and no greater than 2.27%. The chart also indicates the standard deviation, 0.19%, of estimated default rate.

Exhibit 1

Frequency Distribution of Baa 5-year Cumulative Default Rates in 10,000 Bootstrapped Samples

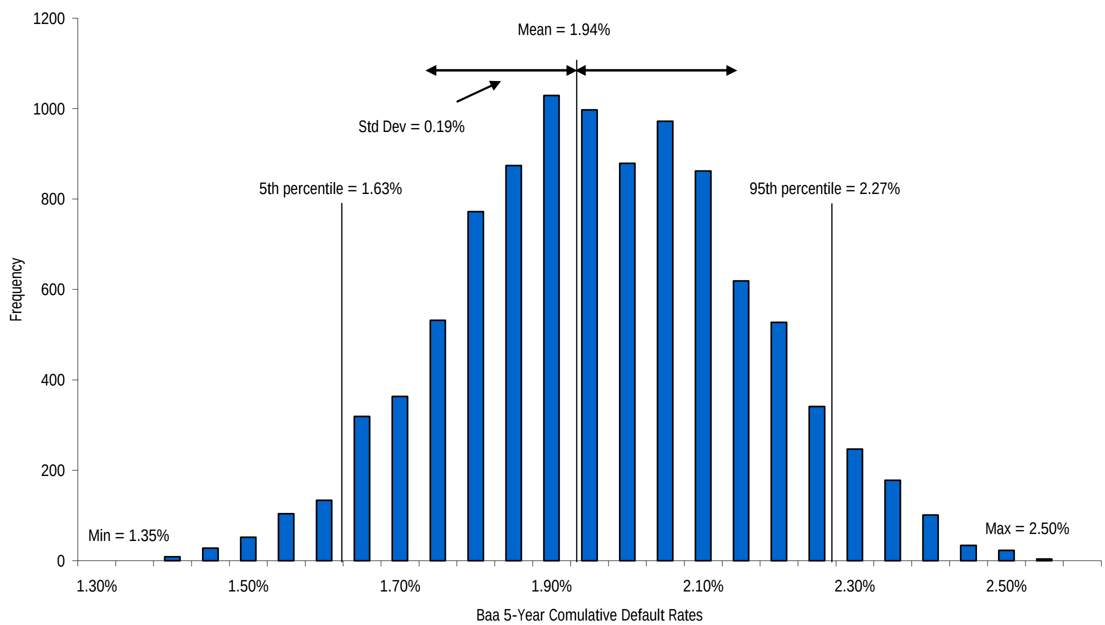

Exhibit 2 presents these sample statistics for the different rating categories (Aaa through Caa) over different time horizons (one through ten years). The table indicates that most rating categories at most horizons have statistically significantly different estimated mean default rates. Some of the highest rating categories (Aaa, Aa, and single A — the low default portfolio ratings in Basel terminology) are not statistically distinguishable at the one-year horizon. They are, however, clearly distinguishable at the two-year horizon, as the confidence interval range the 5th to the 95th percentile ranges from 0.000% to 0.000% for Aaa, from 0.003% for 0.044% for Aa and from 0.061 to 0.144% for single A. All of Moody's long-term ratings would therefore meet the Basel II back-testing criteria for internal rating systems [Basel Committee (2005) page 5].

5. *Moreover, our data sample is almost twice as large; therefore, our confidence intervals for the one-year default rates are considerably tighter than their bootstrap results.*

Exhibit 16

Exhibit 2

Historical Average Cumulative Default Rates and Bootstrapped Confidence Intervals
Derived from Monthly Cohorts: 1970-2006

	Time Horizon (Years) 1	2	3	4	5	6	7	8	9	10
Mean of the Bootstrapped Samples (μ)										
Aaa	0.000%	0.000%	0.000%	0.026%	0.099%	0.172%	0.250%	0.334%	0.424%	0.520%
Aa	0.008%	0.019%	0.042%	0.106%	0.177%	0.260%	0.343%	0.415%	0.463%	0.522%
A	0.021%	0.095%	0.220%	0.344%	0.472%	0.614%	0.759%	0.925%	1.105%	1.286%
Baa	0.181%	0.506%	0.929%	1.433%	1.937%	2.449%	2.956%	3.448%	4.013%	4.633%
Ba	1.203%	3.222%	5.568%	7.953%	10.207%	12.226%	13.992%	15.690%	17.371%	19.095%
B	5.235%	11.298%	17.044%	22.054%	26.791%	30.976%	34.762%	37.972%	40.908%	43.322%
Caa-C	19.466%	30.509%	39.731%	46.935%	52.659%	56.841%	59.965%	63.289%	66.359%	69.251%
Standard Deviation of the Bootstrapped Samples (σ)										
Aaa	n/a	n/a	n/a	0.027%	0.070%	0.104%	0.152%	0.207%	0.268%	0.335%
Aa	0.006%	0.013%	0.022%	0.038%	0.059%	0.084%	0.112%	0.139%	0.159%	0.178%
A	0.008%	0.023%	0.043%	0.061%	0.081%	0.105%	0.129%	0.156%	0.184%	0.214%
Baa	0.026%	0.062%	0.101%	0.145%	0.191%	0.237%	0.287%	0.338%	0.396%	0.462%
Ba	0.080%	0.186%	0.302%	0.417%	0.531%	0.635%	0.731%	0.831%	0.935%	1.044%
B	0.184%	0.376%	0.556%	0.727%	0.903%	1.070%	1.240%	1.409%	1.601%	1.791%
Caa-C	0.747%	1.179%	1.590%	1.959%	2.333%	2.691%	3.044%	3.501%	3.845%	4.616%
Precision of the Bootstrapped Samples (σ/μ)										
Aaa	n/a	n/a	n/a	105.166%	70.155%	60.360%	60.677%	62.056%	63.340%	64.398%
Aa	71.191%	71.383%	51.005%	36.219%	33.521%	32.473%	32.747%	33.513%	34.262%	34.170%
A	38.272%	23.702%	19.502%	17.722%	17.229%	17.084%	17.028%	16.818%	16.689%	16.633%
Baa	14.238%	12.326%	10.917%	10.155%	9.853%	9.685%	9.694%	9.805%	9.863%	9.966%
Ba	6.636%	5.762%	5.415%	5.247%	5.207%	5.196%	5.227%	5.298%	5.380%	5.465%
B	3.511%	3.325%	3.264%	3.296%	3.370%	3.455%	3.567%	3.710%	3.914%	4.135%
Caa-C	3.838%	3.864%	4.002%	4.173%	4.430%	4.734%	5.076%	5.532%	5.793%	6.665%
5th percentile										
Aaa	0.000%	0.000%	0.000%	0.000%	0.010%	0.042%	0.051%	0.051%	0.051%	0.051%
Aa	0.000%	0.003%	0.010%	0.048%	0.085%	0.128%	0.169%	0.199%	0.218%	0.246%
A	0.009%	0.061%	0.153%	0.246%	0.340%	0.448%	0.599%	0.676%	0.819%	0.955%
Baa	0.140%	0.408%	0.769%	1.198%	1.631%	2.069%	2.497%	2.904%	3.387%	3.911%
Ba	1.073%	2.912%	5.081%	7.263%	9.341%	11.203%	12.848%	14.369%	15.887%	17.441%
B	4.939%	10.685%	16.149%	20.894%	25.342%	29.233%	32.755%	35.681%	38.376%	40.469%
Caa-C*	18.270%	28.603%	37.074%	43.709%	48.005%	51.421%	53.537%	n/a	n/a	n/a
95th percentile										
Aaa	0.000%	0.000%	0.000%	0.080%	0.233%	0.362%	0.529%	0.715%	0.911%	1.135%
Aa	0.018%	0.044%	0.082%	0.173%	0.279%	0.410%	0.546%	0.666%	0.745%	0.836%
A	0.035%	0.134%	0.293%	0.449%	0.613%	0.794%	0.982%	1.195%	1.418%	1.648%
Baa	0.225%	0.614%	1.101%	1.682%	2.268%	2.854%	3.440%	4.005%	4.672%	5.413%
Ba	1.332%	3.520%	6.052%	8.637%	11.087%	13.299%	15.218%	17.078%	18.934%	20.848%
B	5.546%	11.940%	17.988%	23.271%	28.304%	32.795%	36.810%	40.348%	43.586%	46.333%
Caa-C*	20.712%	32.590%	42.531%	50.364%	57.533%	62.300%	66.256%	n/a	n/a	n/a
Minimum										
Aaa	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
Aa	0.000%	0.000%	0.000%	0.008%	0.020%	0.043%	0.072%	0.076%	0.076%	0.076%
A	0.001%	0.026%	0.086%	0.159%	0.237%	0.325%	0.383%	0.468%	0.547%	0.638%
Baa	0.101%	0.313%	0.597%	0.955%	1.351%	1.764%	2.068%	2.364%	2.795%	3.274%
Ba	0.886%	2.550%	4.550%	6.606%	8.430%	10.104%	11.368%	12.601%	13.811%	14.951%
B	4.626%	10.174%	15.492%	19.888%	24.075%	27.627%	30.704%	33.560%	35.716%	37.208%
Caa-C	16.842%	26.266%	34.600%	40.674%	45.230%	48.017%	50.627%	51.952%	53.108%	53.108%
Maximum										
Aaa	0.000%	0.000%	0.000%	0.180%	0.499%	0.740%	0.999%	1.277%	1.587%	1.972%
Aa	0.034%	0.076%	0.134%	0.309%	0.472%	0.618%	0.790%	0.996%	1.152%	1.304%
A	0.049%	0.175%	0.379%	0.588%	0.790%	1.011%	1.230%	1.477%	1.726%	2.060%
Baa	0.281%	0.703%	1.255%	1.879%	2.503%	3.158%	3.890%	4.613%	5.377%	6.236%
Ba	1.554%	4.005%	6.814%	9.652%	12.336%	14.646%	16.795%	18.813%	20.722%	22.644%
B	5.899%	12.614%	18.968%	24.534%	29.823%	34.464%	38.632%	42.519%	46.297%	49.932%
Caa-C	22.917%	34.810%	45.222%	54.179%	61.473%	66.251%	70.046%	74.552%	80.095%	86.366%

Note: Data are drawn from the new SRT. The cumulative default rates are created using the monthly cohort method, using data from 1970-2006. The statistics are derived from 10,000 bootstrapped samples.
** The confidence intervals for the Caa-C class are adjusted depending on the percentage of monthly cohorts that are missing from each time horizon.*

Exhibit 16

Exhibit 3 plots the ratios of the standard deviations to the means for these CDRs (the coefficients of variation), which are inversely related to the precision with which these means are estimated. As seen in the chart, for the Baa and the speculative-grade rating categories, the standard error is generally small, at 10% or less than the mean. The mean default rates for the Aaa, Aa, and A rating categories, however, tend to be less precisely estimated, with standard errors ranging from 15% to about 100% of their respective means. The chart shows that means are more precisely estimated for the higher risk rating categories and at longer time horizons. These results are consistent with the well known finding that accurate default rate estimation requires larger datasets when the default rates are lower; i.e., at higher rating categories and at shorter horizons.

Exhibit 3

Ratio of the Standard Deviation to the Mean Cumulative Default Rates for the Bootstrapped Sample

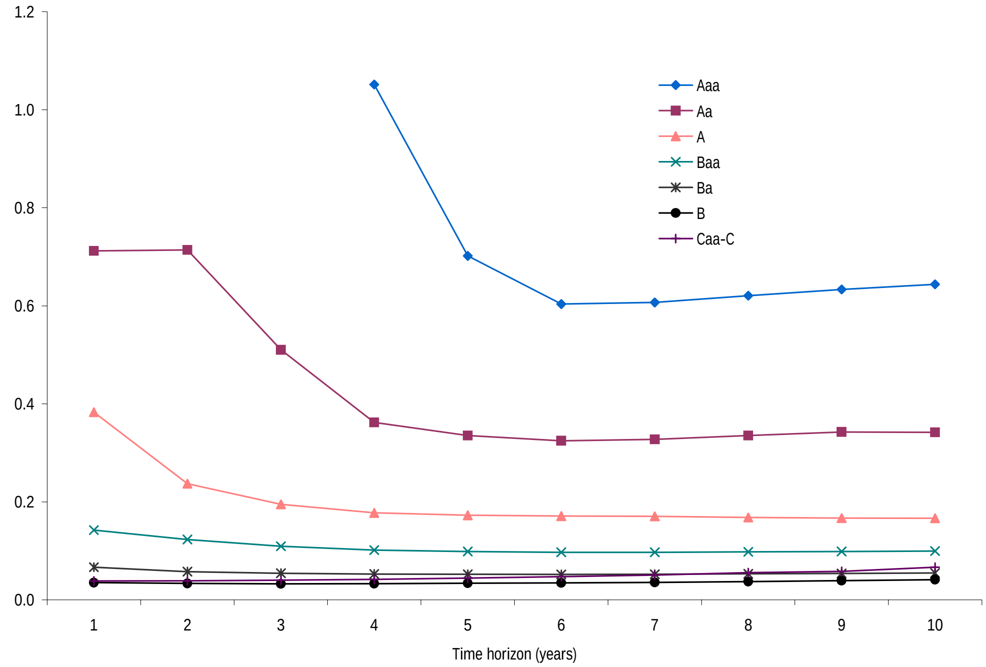

It is worth noting that Moody's cumulative default rates are adjusted for data censoring: i.e. defaults and rating withdrawals that occur prior to the current time interval.[6] Hence, the effective sample size (and the size of the denominators of the marginal default rates) falls as the time horizon lengthens. The decrease in the effective sample size due to adjustments for data censoring is often misunderstood as inducing an upward bias in average cumulative default rate estimates.[7] This view confuses concerns about sample size and statistical significance with bias.

Interestingly, the precision of the default rate estimates does not deteriorate at long time horizons, such as ten years, even though long-term default rates are cumulated from one- year marginal default rates at each horizon, and in some of the cohorts, the number of observations in the tenth year is quite small.[8] The small samples in some cohorts' later years do not cause a problem for the overall sample estimates, because the overall estimates are derived from averages across a large number of cohorts.[9]

Exhibits 2 and 3 reveal that while the standard deviations of default rates by rating category do indeed increase as time horizon lengthens, the precision of the estimates of the mean default rates actually increases with the length of the horizon.

6. *See Cantor and Hamilton (2006).*
7. *See, for example, DeRosa-Farag, S., Blau, J., Matousek, P., Chandra, I. (1999).*
8. *For example, the cohort of B-rated issuers formed on January 1, 1996 contained 519 firms. Ten years later, there were just 84 surviving issuers left from this cohort.*
9. *In the Appendix, we show as expected that the bootstrap estimates of marginal default rates are less precise as the time horizons is lengthened, particularly for speculative-grade companies. (The marginal default rate in year t is the probability of defaulting in year t conditional on having not defaulted through year t-1; whereas, the cumulative default rate in year t is the probability of defaulting during any year before or during year t.)*

Exhibit 16

Bootstrap Estimates of Marginal Default Rates and Their Confidence Intervals

Cumulative default rates are, of course, measured more precisely than marginal default rates (i.e., hazard rates). That is, the estimate of the probability of default over five years is unsurprisingly more precisely estimated that the probability of default in the fifth year, conditional on having survived until the fourth year.

Exhibits 4 and 5 present the same sample statistics for marginal default rates and plots their coefficients of variation (the inverse of the precision of the mean estimate). Interestingly, as the horizon is lengthens, the precision of the estimated mean marginal default rates (as measured by the inverse of the coefficient of variation) improves more slowly than does the precision of the cumulative default rate and, in fact, decreases in many cases.

Exhibit 4

Historical Average Marginal Default Rates and Bootstrapped Confidence Intervals
Derived from Monthly Cohorts: 1970-2006

	Time Horizon (Years)									
	1	2	3	4	5	6	7	8	9	10
Issuer Weighted Average (μ)										
Aaa	0.000%	0.000%	0.000%	0.027%	0.075%	0.073%	0.079%	0.084%	0.090%	0.097%
Aa	0.008%	0.011%	0.023%	0.064%	0.071%	0.082%	0.083%	0.073%	0.048%	0.059%
A	0.021%	0.075%	0.126%	0.124%	0.129%	0.143%	0.146%	0.167%	0.182%	0.183%
Baa	0.182%	0.327%	0.426%	0.510%	0.514%	0.524%	0.521%	0.509%	0.588%	0.649%
Ba	1.202%	2.041%	2.425%	2.529%	2.453%	2.256%	2.020%	1.984%	2.001%	2.094%
B	5.240%	6.408%	6.489%	6.050%	6.085%	5.718%	5.484%	4.919%	4.743%	4.097%
Caa-C	19.488%	13.713%	13.271%	11.988%	10.840%	8.904%	7.321%	8.426%	8.446%	9.069%
Standard Deviation (σ)										
Aaa	0.000%	0.000%	0.000%	0.022%	0.038%	0.044%	0.047%	0.051%	0.055%	0.060%
Aa	0.005%	0.008%	0.013%	0.023%	0.026%	0.031%	0.033%	0.032%	0.027%	0.033%
A	0.008%	0.016%	0.023%	0.024%	0.025%	0.028%	0.030%	0.034%	0.037%	0.038%
Baa	0.025%	0.039%	0.048%	0.055%	0.061%	0.068%	0.069%	0.072%	0.083%	0.098%
Ba	0.084%	0.121%	0.149%	0.163%	0.172%	0.183%	0.196%	0.207%	0.223%	0.248%
B	0.187%	0.226%	0.260%	0.304%	0.369%	0.424%	0.491%	0.515%	0.599%	0.706%
Caa-C	0.724%	0.736%	0.930%	1.174%	1.455%	1.602%	2.063%	2.578%	3.825%	5.555%
Coefficient of Variation (σ/μ)										
Aaa				0.808	0.505	0.600	0.603	0.606	0.609	0.611
Aa	0.687	0.751	0.541	0.368	0.367	0.374	0.400	0.433	0.561	0.555
A	0.370	0.210	0.181	0.194	0.197	0.197	0.204	0.201	0.201	0.207
Baa	0.139	0.121	0.112	0.107	0.119	0.129	0.132	0.142	0.141	0.151
Ba	0.070	0.059	0.061	0.064	0.070	0.081	0.097	0.104	0.111	0.118
B	0.036	0.035	0.040	0.050	0.061	0.074	0.090	0.105	0.126	0.172
Caa-C	0.037	0.054	0.070	0.098	0.134	0.180	0.282	0.306	0.453	0.613
5th percentile										
Aaa	0.000%	0.000%	0.000%	0.000%	0.010%	0.000%	0.000%	0.000%	0.000%	0.000%
Aa	0.000%	0.000%	0.003%	0.027%	0.031%	0.034%	0.028%	0.023%	0.012%	0.000%
A	0.009%	0.048%	0.090%	0.087%	0.090%	0.098%	0.100%	0.112%	0.122%	0.118%
Baa	0.140%	0.263%	0.350%	0.418%	0.417%	0.422%	0.411%	0.393%	0.454%	0.496%
Ba	1.073%	1.843%	2.186%	2.260%	2.178%	1.978%	1.725%	1.664%	1.641%	1.693%
B	4.939%	6.021%	6.049%	5.566%	5.512%	5.077%	4.740%	4.116%	3.753%	3.028%
Caa-C*	18.270%	12.418%	11.592%	9.923%	7.828%	5.452%	3.351%	n/a	n/a	n/a
95th percentile										
Aaa	0.000%	0.000%	0.000%	0.080%	0.161%	0.167%	0.180%	0.193%	0.208%	0.225%
Aa	0.018%	0.027%	0.047%	0.105%	0.120%	0.138%	0.143%	0.129%	0.097%	0.118%
A	0.035%	0.103%	0.164%	0.164%	0.172%	0.192%	0.197%	0.226%	0.248%	0.250%
Baa	0.225%	0.395%	0.503%	0.602%	0.614%	0.628%	0.638%	0.633%	0.732%	0.805%
Ba	1.332%	2.242%	2.660%	2.799%	2.751%	2.553%	2.326%	2.312%	2.353%	2.497%
B	5.546%	6.799%	6.940%	6.560%	6.681%	6.389%	6.293%	5.797%	5.804%	5.241%
Caa-C*	20.712%	15.122%	15.021%	14.146%	14.174%	12.812%	12.052%	n/a	n/a	n/a
Minimum										
Aaa	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
Aa	0.000%	0.000%	0.000%	0.000%	0.000%	0.009%	0.000%	0.000%	0.000%	0.000%
A	0.001%	0.021%	0.057%	0.053%	0.056%	0.047%	0.057%	0.069%	0.075%	0.072%
Baa	0.101%	0.194%	0.280%	0.344%	0.332%	0.321%	0.309%	0.296%	0.291%	0.348%
Ba	0.886%	1.620%	1.901%	2.033%	1.860%	1.604%	1.341%	1.354%	1.245%	1.323%
B	4.626%	5.671%	5.692%	5.071%	5.006%	4.369%	3.825%	3.202%	2.655%	2.045%
Caa-C	16.842%	11.333%	10.476%	8.436%	6.396%	3.829%	1.297%	0.973%	0.000%	0.000%

Exhibit 4

Historical Average Marginal Default Rates and Bootstrapped Confidence Intervals
Derived from Monthly Cohorts: 1970-2006

	Time Horizon (Years)									
	1	2	3	4	5	6	7	8	9	10
Maximum										
Aaa	0.000%	0.000%	0.000%	0.180%	0.355%	0.299%	0.320%	0.341%	0.364%	0.391%
Aa	0.034%	0.047%	0.087%	0.194%	0.182%	0.215%	0.217%	0.208%	0.169%	0.232%
A	0.049%	0.137%	0.219%	0.164%	0.219%	0.249%	0.262%	0.314%	0.343%	0.352%
Baa	0.281%	0.462%	0.576%	0.710%	0.716%	0.772%	0.797%	0.802%	0.909%	1.071%
Ba	1.554%	2.566%	3.004%	3.126%	3.089%	2.898%	2.592%	2.684%	2.830%	3.024%
B	5.899%	7.261%	7.462%	7.164%	7.301%	7.229%	7.017%	6.723%	7.862%	6.770%
Caa-C	22.917%	16.534%	17.398%	16.794%	16.835%	15.936%	16.170%	20.557%	29.469%	45.802%

Note: Data are drawn from the new SRT. The marginal default rates are created using the monthly cohort method, using data from 1970-2006. The statistics are derived from 10,000 bootstrapped samples.
** The confidence intervals for the Caa-C class are adjusted depending on the percentage of monthly cohorts that are missing from each time horizon.*

Exhibit 5

Ratio of the Standard Deviation to the Mean Marginal Default Rates for the Bootstrapped Sample

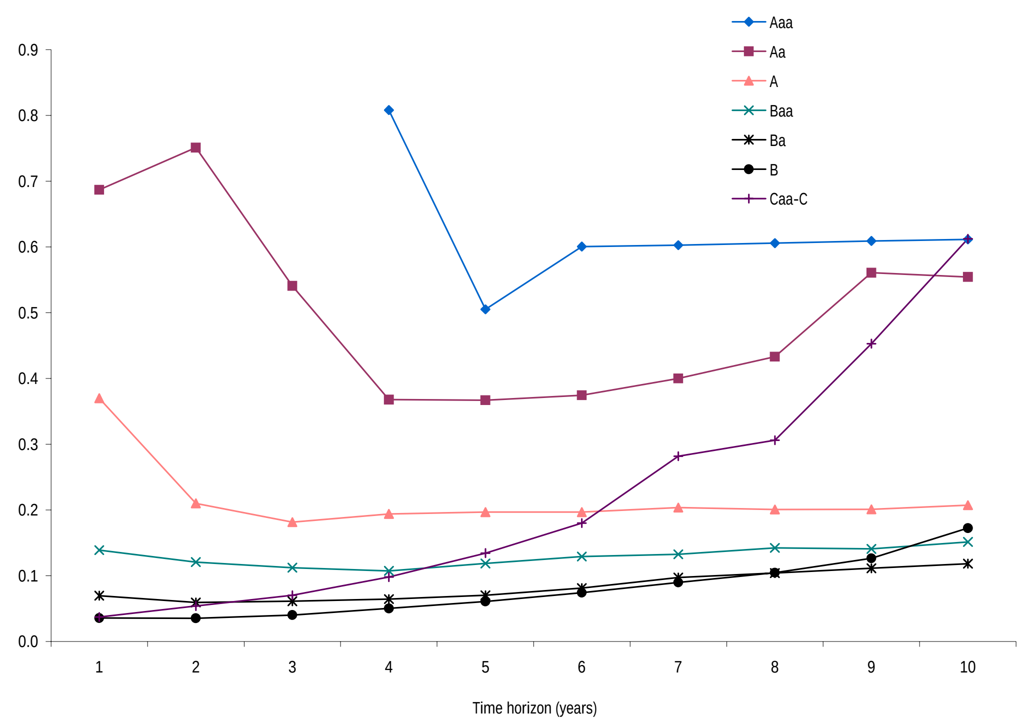

Exhibit 16

Uncertainty around Long-Run Average CDRs for Individual Cohorts

In the previous section, we presented estimates of uncertainty about the long-run average mean. It should be obvious, however, that these measures grossly understate uncertainty about the performance of any *individual* rating cohort. The long-run average default rates smooth out variations in macroeconomic conditions; whereas, the measured performance of any one cohort will reflect the unique macroeconomic conditions that prevailed during its "lifetime."

The high variability of cohort default rates should not be surprising because the primary intent of Moody's ratings management system is to achieve a powerful rank ordering of credit risk within the *current* cohort, and the secondary intent is to maintain rating stability when possible.[10] Moody's does not shift all its ratings up or down with changes in macroeconomic conditions, and, hence, sharp variations in default rates by rating category should be expected over time.

Exhibits 6 and 7 present the historical realizations of five-year default rates for investment-grade and speculative-grade cohorts formed at different monthly intervals.[11] As is evident, default rates vary over time; however, to a very large extent, the expected rank ordering of risk is maintained across rating categories for each monthly cohort.[12]

Exhibit 6

5-Yr Cumulative Default Rates for Monthly Investment-Grade Cohorts

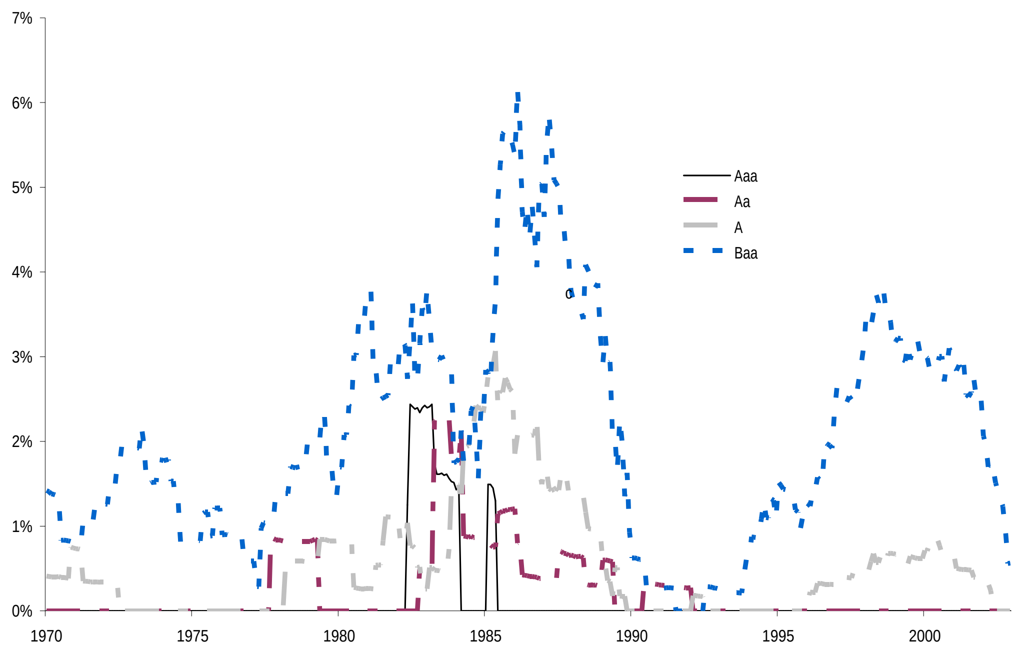

10. *See Cantor and Mann (2006).*
11. *For ease of presentation, we have omitted the Caa default rates from Exhibit 6 which exhibits high volatility at some horizons, for some cohorts, because sample sizes were very small.*
12. *It is important to note that the historical variation in cohort CDRs has two causes, one being changes in macroeconomic conditions (the "underlying" default rate) and the other being random chance due to the finite sample size of any individual cohort. When the samples are large, which is common in the later years of our data sample, most of the default rate fluctuations are clearly attributable to changes in the underlying risk of default rather than chance realizations within the sample. Separating these two effects formally, however, would require adding more structure to the analysis.*

Exhibit 16

Exhibit 7

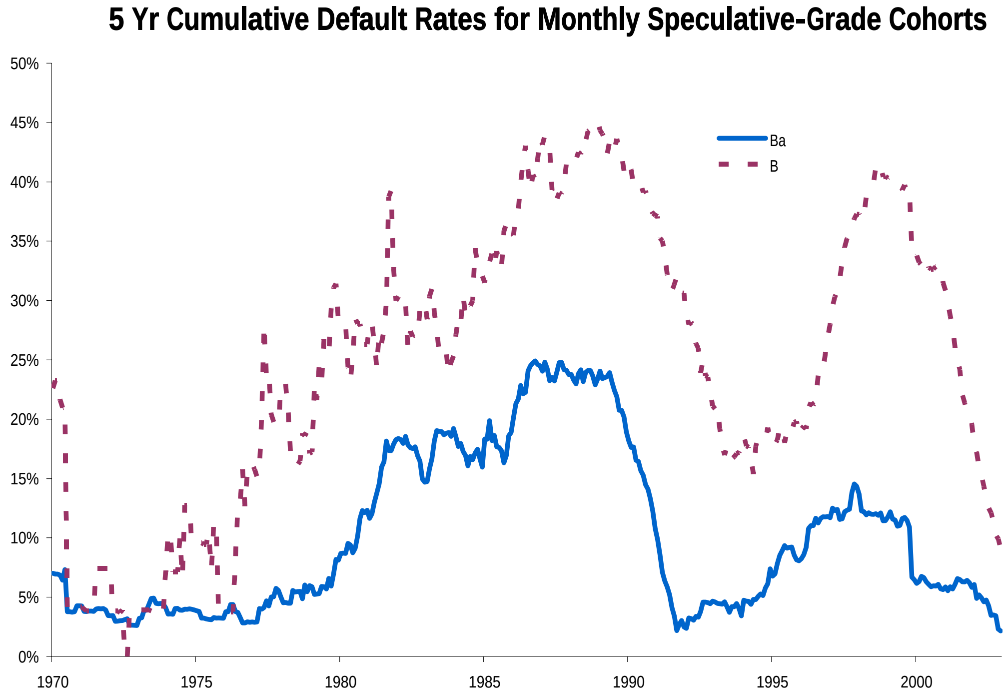

Although the variation in realized default rates across cohorts formed at different points in time has been very large, Moody's ratings have provided consistent and powerful rank orderings of default risk within individual cohorts. For example, since 1992, there has not been a single instance in which a higher-rated cohort has experienced a higher 5-year CDR than a lower-rated cohort that was formed on the same month.

Conclusion

Default rates are key parameters used in a variety of credit risk management and rating applications. In addition to estimates of long-run average default rates for a rating class, estimates of the standard error of the estimates are also important. In this Special Comment we have used bootstrap methods to derive such standard error for the mean cumulative and marginal defaults associated with different rating categories and different horizons. The bootstrap method is a straightforward (though computationally intensive) method for measuring this type of uncertainty.

Our analysis of Moody's data reveals several interesting findings:

- Bootstrap estimates of the standard errors of mean cumulative default rates are relatively tight for lower rating categories and at longer horizons.
- Even for the low-default portfolio portion of the rating scale (Aaa, Aa, and single A), the estimated mean default rates are statistically significantly different from one another at the two-year horizon.
- While our measures of the long-run average cumulative default rates are in many cases fairly precisely estimated, we should not expect that any individual cohort will necessarily perform very close to the mean.

Exhibit 16

Appendix

	-- Number of Issuers			--	Number of Months
Original Rating	**Total**	**Defaulting**	**Ratings Withdrawn**	**Censored at Sample End**	**Average Before Exiting Sample Length**
Aaa	466	2	205	259	175
Aa	1441	13	576	852	119
A	2367	45	915	1407	129
Baa	1944	96	907	941	119
Ba	2002	361	1152	489	84
B	2690	601	1294	795	50
Caa-C	460	131	174	155	34
Total	**11370**	**1249**	**5223**	**4898**	

References

Basel Committee on Banking Supervision (2005), "Validation of Low-Default Portfolios in the Basel II Framework," *Basel Committee Newsletter*, No. 6, September.

Cantor, Richard and Eric Falkenstein (2001), "Testing for Rating Consistency in Annual Default Rates," *Moody's Special Comment*, February, reprinted in the *Journal of Fixed Income*, September.

Cantor, Richard and Christopher Mann (2006), Analyzing the Tradeoff between Ratings Accuracy and Stability," *Moody's Special Comment*, September, reprinted in the *Journal of Fixed Income*, forthcoming, 2007.

Christensen, Jens, Ernst Hansen, and David Lando (2004), "Confidence Sets for Continuous-Time Rating Transition Probabilities," *Journal of Banking & Finance* 28: 2575–2602.

DeRosa-Farag, S., Blau, J., Matousek, P., Chandra, I. (1999), "Default Rates in the High Yield Market," *Journal of Fixed Income*, June: 7-31.

Hamilton, David and Richard Cantor (2004), Rating Transitions and Defaults Conditional on Watchlist, Outlook and Rating History," *Moody's Special Comment*, February, reprinted as "Rating Transition and Default Rates Conditioned on Outlooks," *Journal of Fixed Income*, September.

Hamilton, David and Richard Cantor (2004), "Measuring Corporate Default Rates," *Moody's Special Comment*, November.

Hanson, Samuel and Til Schuermann (2006), "Confidence Intervals for Probabilities of Default," *Journal of Banking & Finance* 30: 2281–2301.

Stein, Roger (2006) "Are the Probabilities Right? Dependent Defaults and the Number of Observations Required to Test for Default Rate Accuracy," *Journal of Investment Management*, 1st Quarter.

Witt, Gary (2004). "Moody's Correlated Binomial Default Distribution," *Moody's Structured Finance Rating Methodology*, August.

Exhibit 16

To order reprints of this report (100 copies minimum), please call 1.212.553.1658.
Report Number: 102662

Author

Richard Cantor

Production Specialist

Yung Louie